

07-03-00  
FISH & RICHARDSON P.C.

June 30, 2000

Attorney Docket No.: 10559/228001/P8793

**Box Patent Application**  
Commissioner for Patents  
Washington, DC 20231

Presented for filing is a new original patent application of:

Applicant: SATYENDRA YADAV, KAPIL JAIN AND ANAND RANGARAIAN

Title: PACKET PROCESSING IN A ROUTER ARCHITECTURE

Enclosed are the following papers, including those required to receive a filing date under 37 CFR 1.53(b):

	Pages
Specification	8
Claims	6
Abstract	1
Declaration	[To be Filed at a Later Date]
Drawing(s)	2

Enclosures:  
— Postcard.

There are 30 total claims, 4 of which are independent.

Basic filing fee	\$0
Total claims in excess of 20 times \$18	\$0
Independent claims in excess of 3 times \$78	\$0
Fee for multiple dependent claims	\$0
Total filing fee:	\$0

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL558600466US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D C. 20231.

Date of Deposit June 30, 2000

Signature *Derek W. Norwood*

Derek W. Norwood  
Typed or Printed Name of Person Signing Certificate

4350 La Jolla Village Drive  
Suite 500  
San Diego, California  
92122

Telephone  
858 678-5070

Facsimile  
858 678-5099

Web Site  
www.fr.com

Frederick P. Fish  
1855-1930

W.K. Richardson  
1859-1951



BOSTON

DALLAS

DELAWARE

NEW YORK

SAN DIEGO

SILICON VALLEY

TWIN CITIES

WASHINGTON, DC

FISH & RICHARDSON P.C.

Commissioner for Patents

June 30, 2000

Page 2

No filing fee is being paid at this time. Please apply any other required fees, **EXCEPT FOR THE FILING FEE**, to deposit account 06-1050, referencing the attorney document number shown above. A duplicate copy of this transmittal letter is attached.

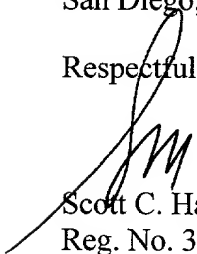
If this application is found to be incomplete, or if a telephone conference would otherwise be helpful, please call the undersigned at (858) 678-5070.

Kindly acknowledge receipt of this application by returning the enclosed postcard.

Please send all correspondence to:

SCOTT C. HARRIS  
Fish & Richardson P.C.  
Customer Number: 20985  
4350 La Jolla Village Drive, Suite 500  
San Diego, CA 92122

Respectfully submitted,



Scott C. Harris  
Reg. No. 32,030  
Enclosures  
SCH/rpi  
10041778.doc

APPLICATION  
FOR  
UNITED STATES LETTERS PATENT

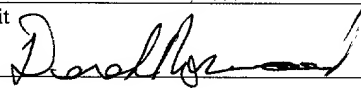
TITLE: PACKET PROCESSING IN A ROUTER ARCHITECTURE  
APPLICANT: SATYENDRA YADAV, KAPIL JAIN AND ANAND  
RANGARAIAN

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL558600466US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit June 30, 2000

Signature 

Derek W. Norwood  
Typed or Printed Name of Person Signing Certificate

## PACKET PROCESSING IN A ROUTER ARCHITECTURE

### BACKGROUND

This invention relates to packet processing in a router architecture.

5       A conventional single-box router processes packets that need to be switched (or forwarded) as well as packets that contain control information called control packets for computing routing tables. A routing table keeps track of routes to particular network destinations. Packet switching  
10 includes the sending of data information in packets, through a network, to a remote location. Each packet of data information has a unique source address and carries its own destination address.

15       A conventional router has a device-driver that controls network interface cards located in the router. The device-driver handles communication between the interface cards between remotely located devices and a networking stack in the router. The networking stack is a stack of software layers such as IP (internet protocol), TCP/UDP (transmission control  
20 protocol/user datagram protocol), and socket library (application programming interface library) between networking applications and the device-driver. A packet traverses the

networking stack to be delivered to a networking application, which handles the processing of control packets of a networking device.

5

## DESCRIPTION OF DRAWINGS

FIGS. 1 and 2 illustrate receiving and transmitting a packet in a distributed router architecture according to an embodiment of the invention.

10

## DETAILED DESCRIPTION

15

A router architecture described below separates the processing of packet switching (or forwarding) from the processing of control packets. FIG. 1 shows an implementation of a distributed router architecture 1000 in which switches 100, 200 for handling packet switching and a controller 500 for processing control packets are distributed across a network such as an Ethernet network 50. The controller 500 controls the switches 100, 200.

20

The distributed router architecture 1000 may be made compatible with networking applications 502, such as telnet, that work in a conventional router. One way is by keeping the distributed switches 100, 200 as hidden from the networking applications 502. Otherwise, the networking applications 502

would require modification to work under the distributed architecture. Hiding the fact that the switches 100, 200 are distributed can provide a single router view of the switches 100, 200 to the networking applications 502. From the  
5 perspective of the networking applications 502, packets may appear to come from a router and not through a separate switch.

A known single-box router contains network interfaces, each controlled by a device-driver. In the distributed router  
10 architecture 1000, the network interfaces that communicate with other routers or hosts are placed in the switches 100, 200, but not in the controller 500. FIG. 1 shows network interfaces A, B, and C located in the switch 100 and network interfaces D and E located in the switch 200. To "hide" the  
15 fact that the switches are distributed over the Ethernet 50, a pseudo device-driver 512 can be introduced in the controller 500. The pseudo device-driver 512 causes control packets, which are packets containing routing information for computing routing tables, to be delivered to the networking stack (e.g.  
20 socket library 504, TCP/UDP 506, and IP 510) of the controller 500 as if coming from virtual interfaces Av, Bv, Cv, Dv, or Ev in the pseudo device-driver 512. Each virtual interface corresponds to an interface present in a switch. For example, the virtual interface Av in the controller 500 corresponds to

the interface A in the switch 100, and so on. To the networking applications 502, the controller 500 appears to be a conventional single-box router with as many interfaces as there are in all the switches.

5           A control packet that is received by one of switch's interfaces can be relayed to the controller 500 and appear as input to its corresponding virtual interface. Conversely, a control packet sent out through one of the virtual interfaces in the controller 500 can be relayed to an appropriate switch  
10           containing the corresponding interface. The receiving and sending of packets as described above can be accomplished using an IP encapsulation protocol termed a virtual interface protocol (VIP). The VIP is implemented by a VIP module 508 located in the controller 500 and VIP modules 104 and 204  
15           located in the switches 100 and 200, respectively. The encapsulation carries frames of one protocol which includes as the data in another protocol.

FIG. 1 illustrates how a control packet is received by the switch 100 that appears to reside in the controller 500  
20           and delivered to a networking application 502. In phase 1, the switch 100 receives a packet through the interface C. In phase 2, the switch 100 appends an external IP header to the packet and delivers the packet to the IP module 106. A protocol field in the external IP header is filled with

information regarding the VIP. The external header provides the destination address of the controller 500. In phase 3, the IP module 106, which has access to a routing table, performs route lookup for the controller 500 and sends the packet through an interface 102 specified in the routing table. In phase 4, the packet reaches the controller 500 at an interface 514 from the interface 102 and is delivered to the interface CV in the pseudo device driver 512. In phase 5, the interface Cv in the pseudo device-driver 512 delivers the packet to the IP module 510. In phase 6, the IP module 510 delivers the packet to a VIP module 508 based on the protocol field in the external IP header. In phase 7, the VIP module 508 strips off the external IP header and delivers the packet to the IP module 510. The packet is delivered to the IP module 510 as if it were coming from the virtual interface Cv corresponding to the interface C in the switch 100. In phases 8 to 10, the packet traverses the networking stack--IP 510, TCP/UDP 506, and Socket library 504--and the data is delivered to the application 502.

FIG. 2 illustrates how a control packet is transmitted from a networking application 502. In phases 11 to 13, a packet is written by the application 502 using the socket library 504 and gets encapsulated by the TCP/UDP module 506. It is then given to the IP module 510. In phase 14, the IP



module 510 appends an IP header and performs route lookup for the destination and sends it through the appropriate interface in the pseudo device-driver specified in the routing table. A conventional single-box router software can be used to compute the routing table for the pseudo or actual interfaces of distributed switches. In FIG. 2, the interface Ev is assumed to be the appropriate virtual interface. In phase 15, the pseudo device-driver 512 appends an external IP header to the packet. The destination address contained in the external IP header is that of the switch 200 that contains the interface E, which corresponds to the virtual interface Ev. The protocol field of the external IP header is filled with information regarding the VIP. The pseudo device-driver 512 performs a route lookup for the destination switch 200 and sends the packet to the switch 200 through the interface 514. In phase 16, the packet reaches the switch 200 at an interface device-driver 202. In phase 17, the interface device-driver 202 delivers the packet to an IP module 206. The IP module 206 in phase 18 sends the packet to a VIP module 204 based on the protocol field in the external IP header. In phase 19, the VIP module 204 strips off the external IP header. It then looks into the internal IP header for a source IP address. This source IP address corresponds to one of the network interfaces, the network interface E. In phase 20, the packet

is sent out through the corresponding interface E to a network interface.

In the foregoing example, the controller 500 and the switches 100, 200 were distributed across the Ethernet 50.

5 However, the controller 500 and the switches 100, 200 can be tightly coupled using a high-speed switching fabric instead of being distributed. The routing architecture and the method described above can also be employed for the tightly coupled controller and the switches.

10 Routing architectures represented by the distributed architecture 1000 and the tightly coupled architecture are scalable, provide a single point of management at the controller, and reduce the load on the network. Additionally, the logic circuit of the switches 100, 200 can be made simpler  
15 and more efficient. The foregoing distributed architecture can be implemented in any Network Operating System (NOS) without any changes to its networking stack and routing infrastructure such as routing protocols or a routing table manager. Any networking application that works in an NOS will  
20 work without any changes in the exemplary disclosed architecture. The distributed architecture would also be useful in creating a number of applications such as firewalls for distributed switches with the feature of single point management.

The foregoing techniques can be implemented, for example, in a computer program executable on a computer. The computer program can be stored on a storage medium, such as random access memory (RAM), readable by a general or special purpose programmable computer, for routing packets under the exemplary  
5 disclosed architecture.

Other implementations are within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A routing system comprising:  
a switch configured to switch packets; and  
a controller configured to control the switch and to receive a packet from the switch using a protocol that operates as if the switch were located in the controller.
2. The routing system of claim 1 wherein the switch is arranged to append address information of the controller to the packet for delivery to the controller.
3. The routing system of claim 1 wherein the controller is arranged to strip address information associated with the controller from the packet received from the switch.
4. The routing system of claim 1 wherein the switch has interfaces, each interface corresponding to a virtual interface in the controller.
5. The routing system of claim 4 wherein the packet received by an interface of the switch is sent to the corresponding virtual interface of the controller.

6. The routing system of claim 1 comprising multiple switches and wherein the controller is configured to control the switches.

7. The routing system of claim 1 further comprising a network medium between the switch and controller.

8. The routing system of claim 7 wherein the network medium comprises Ethernet.

9. The routing system of claim 1 wherein the controller comprises a networking stack for receiving and processing the packet.

10. A routing system comprising:  
a switch configured to switch packets; and  
a controller configured to control the switch and to transmit a packet through the switch using a protocol that operates as if the switch were located in the controller.

11. The routing system of claim 10 wherein the controller with the packet is arranged to append address information of the switch to the packet for delivery to the switch.

12. The routing system of claim 11 wherein the switch is arranged to strip the address information associated with the switch from the packet received from the controller.

13. The routing system of claim 10 wherein the switch has interfaces, each interface corresponding to a virtual interface in the controller.

14. The routing system of claim 13 wherein the packet from a virtual interface in the controller is sent to the corresponding interface of the switch.

15. The routing system of claim 10 comprising multiple switches and wherein the controller is configured to control the switches.

16. The routing system of claim 10 wherein the switch and the controller communicate through a network medium.

17. The routing system of claim 16 wherein the network medium comprises Ethernet.

18. The routing system of claim 10 wherein the controller comprises a networking stack for transmitting the packet.

19. A method of routing a packet comprising:  
sending a packet to a first routing component;  
encapsulating the packet with address information;  
sending the encapsulated packet to a second routing  
component based on the address information; and  
unencapsulating the packet at the second component.

20. The method of claim 19 further comprising: sending  
the packet up a networking stack to a networking application.

21. The method of claim 19 wherein the first component  
comprises a switch and the second component comprises a  
controller.

22. The method of claim 21 further comprising:  
looking up a routing table to determine the address of  
the controller before sending the packet to the controller.

23. The method of claim 19 wherein the first component  
comprises a controller and the second component comprises a  
switch.

24. The method of claim 23 further comprising:

looking up a routing table to determine the address of the switch before sending the packet to the switch.

25. An article comprising a computer-readable medium that stores computer-executable instructions for causing a computer system to:

send a packet to a first router component;  
encapsulate the packet with address information;  
send the encapsulated packet to a second router component based on the address information; and  
unencapsulate the packet at the second component.

26. The article of claim 25 further causing the computer system to:

send the packet up a networking stack to a networking application.

27. The article of claim 25 wherein the first component comprise a switch and the second component comprises a controller.

28. The article of claim 27 further causing the computer system to:



look up a routing table to determine the address of the controller before sending the packet to the controller.

29. The article of claim 25 wherein the first component comprises a controller and the second component comprises a switch.

30. The article of claim 27 further causing the computer system to:

look up a routing table to determine the address of the switch before sending the packet to the switch.

ABSTRACT

A routing architecture comprising at least one or more switches for packet switching and a controller for controlling the switches is proposed. A packet is received by the controller through the switch as if the switch is located in the controller.

1000




FIG. 1

1000




FIG. 2